

On Memoryless Quantitative Objectives

Krishnendu Chatterjee¹, Laurent Doyen², and Rohit Singh³

¹ Institute of Science and Technology(IST) Austria

² LSV, ENS Cachan & CNRS, France

³ Indian Institute of Technology(IIT) Bombay

Abstract. In two-player games on graph, the players construct an infinite path through the game graph and get a reward computed by a payoff function over infinite paths. Over weighted graphs, the typical and most studied payoff functions compute the limit-average or the discounted sum of the rewards along the path. Beside their simple definition, these two payoff functions enjoy the property that memoryless optimal strategies always exist.

In an attempt to construct other simple payoff functions, we define a class of payoff functions which compute an (infinite) weighted average of the rewards. This new class contains both the limit-average and discounted sum functions, and we show that they are the only members of this class which induce memoryless optimal strategies, showing that there is essentially no other simple payoff functions.

1 Introduction

Two-player games on graphs have many applications in computer science, such as the synthesis problem [7], and the model-checking of open reactive systems [1]. Games are also fundamental in logics, topology, and automata theory [18, 15, 21]. Games with quantitative objectives have been used to design resource-constrained systems [28, 9, 3, 4], and to support quantitative model-checking and robustness [5, 6, 27].

In a two-player game on a graph, a token is moved by the players along the edges of the graph. The set of states is partitioned into player-1 states from which player 1 moves the token, and player-2 states from which player 2 moves the token. The interaction of the two players results in a play, an infinite path through the game graph. In qualitative zero-sum games, each play is winning for one of the player; in quantitative games, a payoff function assigns a value to every play, which is paid by player 2 to player 1. Therefore, player 1 tries to maximize the payoff while player 2 tries to minimize it. Typically, the edges of the graph carry a reward, and the payoff is computed as a function of the infinite sequences of rewards on the play.

Two payoff functions have received most of the attention in literature: the *mean-payoff* function (for example, see [11, 28, 16, 20, 12, 22]) and the *discounted-sum* function (for example, see [25, 12, 23, 24, 9]). The mean-payoff value is the long-run average of the rewards. The discounted sum is the infinite sum of the rewards under a discount factor $0 < \lambda < 1$. For an infinite sequence of rewards $w = w_0w_1\dots$, we have:

$$\text{MeanPayoff}(w) = \liminf_{n \rightarrow \infty} \frac{1}{n} \cdot \sum_{i=0}^{n-1} w_i \quad \text{DiscSum}_\lambda(w) = (1 - \lambda) \cdot \sum_{i=0}^{\infty} \lambda^i \cdot w_i$$

While these payoff functions have a simple, intuitive, and mathematically elegant definition, it is natural to ask why they are playing such a central role in the study of quantitative games. One answer is perhaps that *memoryless* optimal strategies exist for these objectives. A strategy is memoryless if it is independent of the history of the play and depends only on the current state. Related to this property is the fact that the problem of deciding the winner in such games is in $\text{NP} \cap \text{coNP}$, while no polynomial time algorithm is known for this problem. The situation is similar to the case of parity games in the setting of qualitative games where it was proved that the parity objective is the only prefix-independent objective to admit memoryless winning strategies [8], and the parity condition is known as a canonical way to express ω -regular languages [26].

In this paper, we prove a similar result in the setting of quantitative games. We consider a general class of payoff functions which compute an infinite weighted average of the rewards. The payoff functions are parameterized by an infinite sequence of rational coefficients $\{c_n\}_{n \geq 0}$, and defined as follows:

$$\text{WeightedAvg}(w) = \liminf_{n \rightarrow \infty} \frac{\sum_{i=0}^n c_i \cdot w_i}{\sum_{i=0}^n c_i}.$$

We consider this class of functions for its simple and natural definition, and because it generalizes both mean-payoff and discounted-sum which can be obtained as special cases, namely for $c_i = 1$ for all⁴ $i \geq 0$, and $c_i = \lambda^i$ respectively. We study the problem of characterizing which payoff functions in this class admit memoryless optimal strategies for both players. Our results are as follows:

1. If the series $\sum_{i=0}^{\infty} c_i$ converges (and is finite), then discounted sum is the *only* payoff function that admits memoryless optimal strategies for both players.
2. If the series $\sum_{i=0}^{\infty} c_i$ does not converge, but the sequence $\{c_n\}_{n \geq 0}$ is bounded, then for memoryless optimal strategies the payoff function is equivalent to the mean-payoff function (equivalent for the optimal value and optimal strategies of both players).

Thus our results show that the discounted sum and mean-payoff functions, beside their elegant and intuitive definition, are the only members from a large class of natural payoff functions that are simple (both players have memoryless optimal strategies). In other words, there is essentially no other simple payoff functions in the class of weighted infinite average payoff functions. This further establishes the canonicity of the mean-payoff and discounted-sum functions, and suggests that they should play a central role in the emerging theory of quantitative automata and languages [10, 17, 2, 5].

In the study of games on graphs, characterizing the classes of payoff functions that admit memoryless strategies is a research direction that has been investigated in the works of [13, 14] which give general conditions on the payoff functions such that both players have memoryless optimal strategies, and [19] which presents similar results when only one player has memoryless optimal strategies. The conditions given in these previous works are useful in this paper, in particular the fact that it is sufficient to check

⁴ Note that other sequences also define the mean-payoff function, such as $c_i = 1 + 1/2^i$.

that memoryless strategies are sufficient in one-player games [14]. However, conditions such as sub-mixing and selectiveness of the payoff function are not immediate to establish, especially when the sum of the coefficients $\{c_n\}_{n \geq 0}$ does not converge. We identify the necessary condition of boundedness of the coefficients $\{c_n\}_{n \geq 0}$ to derive the mean-payoff function. Our results show that if the sequence is convergent, then discounted sum (specified as $\{\lambda^n\}_{n \geq 0}$, for $\lambda < 1$) is the only memoryless payoff function; and if the sequence is divergent and bounded, then mean-payoff (specified as $\{\lambda^n\}_{n \geq 0}$ with $\lambda = 1$) is the only memoryless payoff function. However we show that if the sequence is divergent and unbounded, then there exists a sequence $\{\lambda^n\}_{n \geq 0}$, with $\lambda > 1$, that does not induce memoryless optimal strategies.

2 Definitions

Game graphs. A two-player game graph $G = \langle Q, E, w \rangle$ consists of a finite set Q of states partitioned into player-1 states Q_1 and player-2 states Q_2 (i.e., $Q = Q_1 \cup Q_2$), and a set $E \subseteq Q \times Q$ of edges such that for all $q \in Q$, there exists (at least one) $q' \in Q$ such that $(q, q') \in E$. The weight function $w : E \rightarrow \mathbb{Q}$ assigns a reward to each edge. For a state $q \in Q$, we write $E(q) = \{r \in Q \mid (q, r) \in E\}$ for the set of successor states of q . A *player-1 game* is a game graph where $Q_1 = Q$ and $Q_2 = \emptyset$. Player-2 games are defined analogously.

Plays and strategies. A game on G starting from a state $q_0 \in Q$ is played in rounds as follows. If the game is in a player-1 state, then player 1 chooses the successor state from the set of outgoing edges; otherwise the game is in a player-2 state, and player 2 chooses the successor state. The game results in a *play* from q_0 , i.e., an infinite path $\rho = \langle q_0 q_1 \dots \rangle$ such that $(q_i, q_{i+1}) \in E$ for all $i \geq 0$. We write Ω for the set of all plays. The prefix of length n of ρ is denoted by $\rho(n) = q_0 \dots q_n$. A strategy for a player is a recipe that specifies how to extend plays. Formally, a *strategy* for player 1 is a function $\sigma : Q^* Q_1 \rightarrow Q$ such that $(q, \sigma(\rho \cdot q)) \in E$ for all $\rho \in Q^*$ and $q \in Q_1$. The strategies for player 2 are defined analogously. We write Σ and Π for the sets of all strategies for player 1 and player 2, respectively.

An important special class of strategies are *memoryless* strategies which do not depend on the history of a play, but only on the current state. Each memoryless strategy for player 1 can be specified as a function $\sigma : Q_1 \rightarrow Q$ such that $\sigma(q) \in E(q)$ for all $q \in Q_1$, and analogously for memoryless player 2 strategies.

Given a starting state $q \in Q$, the *outcome* of strategies $\sigma \in \Sigma$ for player 1, and $\pi \in \Pi$ for player 2, is the play $\omega(s, \sigma, \pi) = \langle q_0 q_1 \dots \rangle$ such that : $q_0 = q$ and for all $k \geq 0$, if $q_k \in Q_1$, then $\sigma(q_0, q_1, \dots, q_k) = q_{k+1}$, and if $q_k \in Q_2$, then $\pi(q_0, q_1, \dots, q_k) = q_{k+1}$.

Payoff functions, optimal strategies. The objective of player 1 is to construct a play that maximizes a *payoff function* $\phi : \Omega \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ which is a measurable function that assigns to every value a real-valued payoff. The value for player 1 is the maximal payoff that can be achieved against all strategies of the other player. Formally the value for player 1 for a starting state q is defined as

$$val_1(\phi) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \phi(\omega(q, \sigma, \pi)).$$

A strategy σ^* is *optimal* for player 1 from q if the strategy achieves at least the value of the game against all strategies for player 2, i.e.,

$$\inf_{\pi \in \Pi} \phi(\omega(q, \sigma^*, \pi)) = \text{val}_1(\phi).$$

The values and optimal strategies for player 2 are defined analogously.

The mean-payoff and discounted-sum functions are examples of payoff functions that are well studied, probably because they are simple in the sense that they induce memoryless optimal strategies and that this property yields conceptually simple fix-point algorithms for game solving [25, 11, 28, 12]. In an attempt to construct other simple payoff functions, we define the class of *weighted average payoffs* which compute (infinite) weighted averages of the rewards, and we ask which payoff functions in this class induce memoryless optimal strategies.

We say that a sequence $\{c_n\}_{n \geq 0}$ of rational numbers has *no zero partial sum* if $\sum_{i=0}^n c_i \neq 0$ for all $n \geq 0$. Given a sequence $\{c_n\}_{n \geq 0}$ with no zero partial sum, the *weighted average payoff function* for a play $\langle q_0 q_1 q_2 \dots \rangle$ is

$$\phi(q_0 q_1 q_2 \dots) = \liminf_{n \rightarrow \infty} \frac{\sum_{i=0}^n c_i \cdot w(q_i, q_{i+1})}{\sum_{i=0}^n c_i}.$$

Note that we use $\liminf_{n \rightarrow \infty}$ in this definition because the plain limit may not exist in general. The behavior of the weighted average payoff functions crucially depends on whether the series $S = \sum_{i=0}^{\infty} c_i$ converges or not. In particular, the plain limit exists if S converges (and is finite). Accordingly, we consider the cases of converging and diverging sum of weights to characterize the class of weighted average payoff functions that admit memoryless optimal strategies for both players. Note that the case where $c_i = 1$ for all $i \geq 0$ gives the mean-payoff function (and S diverges), and the case $c_i = \lambda^i$ for $0 < \lambda < 1$ gives the discounted sum with discount factor λ (and S converges). All our results hold if we consider $\limsup_{n \rightarrow \infty}$ instead of $\liminf_{n \rightarrow \infty}$ in the definition of weighted average objectives.

In the sequel, we consider payoff functions $\phi : \mathbb{Q}^\omega \rightarrow \mathbb{R}$ with the implicit assumption that the value of a play $q_0 q_1 q_2 \dots \in Q^\omega$ according to ϕ is $\phi(w(q_0, q_1)w(q_1, q_2) \dots)$ since the sequence of rewards determines the payoff value.

We recall the following useful necessary condition for memoryless optimal strategies to exist [14]. A payoff function ϕ is *monotone* if whenever there exists a finite sequence of rewards $x \in \mathbb{Q}^*$ and two sequences $u, v \in \mathbb{Q}^\omega$ such that $\phi(xu) \leq \phi(xv)$, then $\phi(yu) \leq \phi(yv)$ for all finite sequence of rewards $y \in \mathbb{Q}^*$.

Lemma 2.1 ([14]). *If the payoff function ϕ induces memoryless optimal strategy for all two-player game graphs, then ϕ is monotone.*

3 Weighted Average with Converging Sum of Weights

The main result of this section is that for converging sum of weights (i.e., if $\lim_{n \rightarrow \infty} \sum_{i=0}^n c_i = c^* \in \mathbb{R}$), the only weighted average payoff function that induce memoryless optimal strategies is the discounted sum.




Fig. 1. Examples of one-player game graphs.

Theorem 3.1. Let $(c_n)_{n \in \mathbb{N}}$ be a sequence of real numbers with no zero partial sum such that $\sum_{i=0}^{\infty} c_i = c^* \in \mathbb{R}$. The weighted average payoff function defined by $(c_n)_{n \in \mathbb{N}}$ induces optimal memoryless strategies for all 2-player game graphs if and only if there exists $0 \leq \lambda < 1$ such that $c_{i+1} = \lambda \cdot c_i$ for all $i \geq 0$.

To prove Theorem 3.1, we first use its assumptions to obtain necessary conditions for the weighted average payoff function defined by $(c_n)_{n \in \mathbb{N}}$ to induce optimal memoryless strategies. By *assumptions of Theorem 3.1*, we refer to the fact that $(c_n)_{n \in \mathbb{N}}$ is a sequence of real numbers with no zero partial sum such that $\sum_{i=0}^{\infty} c_i = c^* \in \mathbb{R}$, and that it defines a weighted average payoff function that induces optimal memoryless strategies for all 2-player game graphs. All lemmas of this section use the the assumptions of Theorem 3.1, but we generally omit to mention them.

Let $d_n = \sum_{i=0}^{n-1} c_i$, $l = \liminf_{n \rightarrow \infty} \frac{1}{d_n}$ and $L = \limsup_{n \rightarrow \infty} \frac{1}{d_n}$. The assumption that $\sum_{i=0}^{\infty} c_i = c^* \in \mathbb{R}$ implies that $l \neq 0$.

Note that $c_0 \neq 0$ since $(c_n)_{n \in \mathbb{N}}$ is a sequence with no zero partial sum. We can define the sequence $c'_n = \frac{c_n}{c_0}$ which defines the same payoff function ϕ . Therefore we assume without loss of generality that $c_0 = 1$.

Lemma 3.1. If the weighted average payoff function defined by $(c_n)_{n \in \mathbb{N}}$ induces optimal memoryless strategies for all 2-player game graphs, then $0 \leq l \leq L \leq 1$.

Proof. Consider the one-player game graph G_1 shown in Fig. 1. In one-player games, strategies correspond to paths. The two memoryless strategies give the paths 0^ω and 1^ω with payoff value 0 and 1 respectively. The strategy which takes edge with reward 1 once, and then always the edge with reward 0 gets payoff $\phi(10^\omega) = \liminf_{n \rightarrow \infty} \frac{1}{d_n} = l$. Similarly, the path 01^ω has payoff $\phi(01^\omega) = \liminf_{n \rightarrow \infty} \left(1 - \frac{1}{d_n}\right) = 1 - \limsup_{n \rightarrow \infty} \frac{1}{d_n} = 1 - L$. Since all such payoffs must be between the payoffs obtained by the only two memoryless strategies, we have $l \geq 0$ and $L \leq 1$, and the result follows ($L \geq l$ follows from their definition). \square

Lemma 3.2. There exists $w_0 \in \mathbb{N}$ such that $w_0 l > 1$ and the following inequalities hold, for all $k \geq 0$: $c_k l \leq 1 - d_k L$ and $c_k w_0 l \geq 1 - d_k L$.

Proof. Since $1 \geq l > 0$ (by Lemma 3.1), we can choose $w_0 \in \mathbb{N}$ such that $w_0 l > 1$. Consider the game graph G_2 shown in Fig. 1 and the case when $w = 1$. The optimal memoryless strategy is to stay on the starting state forever because $\phi(10^\omega) = l \leq \phi(1^\omega) = 1$. Using Lemma 2.1, we conclude that since $\phi(10^\omega) \leq \phi(1^\omega)$, we must have $\phi(0^k 10^\omega) \leq \phi(0^k 1^\omega)$ i.e. $c_k l \leq 1 - \left(\sum_{i=0}^{k-1} c_i\right) L$ which implies $c_k l \leq 1 - d_k L$.

Now consider the case when $w = w_0$ in Fig. 1. The optimal memoryless strategy is to choose the edge with reward w_0 from the starting state since $\phi(w_0 0^\omega) = w_0 l > \phi(1^\omega) = 1$. Using Lemma 2.1, we conclude that since $\phi(w_0 0^\omega) > \phi(1^\omega)$, we must have $\phi(0^k w_0 0^\omega) \geq \phi(0^k 1^\omega)$ i.e. $c_k w_0 l \geq 1 - \left(\sum_{i=0}^{k-1} c_i\right) L$ which implies $c_k w_0 l \geq 1 - d_k L$. \square

From the inequalities in Lemma 3.2, it is easy to see that since $w_0 > 1$ we must have $c_k \geq 0$ for all k .

Corollary 3.1. *Assuming $c_0 = 1$, we have $c_k \geq 0$ for all $k \geq 0$.*

It follows from Corollary 3.1 that the sequence $(d_n)_{n \geq 0}$ is increasing and bounded from above (if d_n was not bounded, then there would exist a subsequence (d_{n_k}) which diverges, implying that the sequence $\{\frac{1}{d_{n_k}}\}$ converges to 0 in contradiction with the fact that $\liminf_{n \rightarrow \infty} \frac{1}{d_n} = l > 0$). Therefore, d_n must converge to some real number say $c^* > 0$ (since $c_0 = 1$). We need a last lemma to prove Theorem 3.1. Recall that we have $c_i \geq 0$ for all i and $\sum_{i=0}^{\infty} c_i = c^* > 0$. Given a finite game graph G , let W be the largest reward in absolute value. For any sequence of rewards (w_n) in a run on G , the sequence $\chi_n = \sum_{i=0}^n c_i(w_i + W)$ is increasing and bounded from above by $2 \cdot W d_n$ and thus by $2 \cdot W c^*$. Therefore, χ_n is a convergent sequence and $\sum_{i=0}^{\infty} c_i w_i$ converges as well. Now, we can write the payoff function as $\phi(w_0 w_1 \dots) = \frac{\sum_{i=0}^{\infty} c_i w_i}{c^*}$. We decompose c^* into $S_0 = \sum_{i=0}^{\infty} c_{2i}$ and $S_1 = \sum_{i=0}^{\infty} c_{2i+1}$, i.e. $c^* = S_0 + S_1$. Note that S_0 and S_1 are well defined.

Lemma 3.3. *If there exist numbers α, β, γ such that $\alpha S_0 + \beta S_1 \leq \gamma(S_0 + S_1)$, then $(\gamma - \alpha)c_i \geq (\beta - \gamma)c_{i+1}$ for all $i \geq 0$.*

Proof. Consider the game graph G_4 as shown in Fig. 1. The condition $\alpha S_0 + \beta S_1 \leq \gamma(S_0 + S_1)$ implies that the optimal memoryless strategy is to always choose the edge with reward γ . This means that $\phi(\gamma^i \alpha \beta \gamma^\omega) \leq \phi(\gamma^\omega)$ hence $\alpha c_i + \beta c_{i+1} \leq \gamma(c_i + c_{i+1})$, i.e. $(\gamma - \alpha)c_i \geq (\beta - \gamma)c_{i+1}$ for all $i \geq 0$. \square

We are now ready to prove the main theorem of this section.

Proof (of Theorem 3.1). First, we show that $S_1 \leq S_0$. By contradiction, assume that $S_1 > S_0$. Choosing $\alpha = 1$, $\beta = -1$, and $\gamma = 0$ in Lemma 3.3, and since $S_0 - S_1 \leq 0$, we get $-c_i \geq -c_{i+1}$ for all $i \geq 0$ which implies $c_n \geq c_0 = 1$ for all n , which contradicts that $\sum_{i=0}^{\infty} c_i$ converges to $c^* \in \mathbb{R}$.

Now, we have $S_1 \leq S_0$ and let $\lambda = \frac{S_1}{S_0} \leq 1$. Consider a sequence of rational numbers $\frac{l_n}{k_n}$ converging to λ from the right, i.e., $\frac{l_n}{k_n} \geq \lambda$ for all n , and $\lim_{n \rightarrow \infty} \frac{l_n}{k_n} = \lambda$. Taking $\alpha = 1$, $\beta = k_n + l_n + 1$, and $\gamma = l_n + 1$ in Lemma 3.3, and since the condition $S_0 + (k_n + l_n + 1)S_1 \leq (l_n + 1)(S_0 + S_1)$ is equivalent to $k_n S_1 \leq l_n S_0$ which holds since $\frac{l_n}{k_n} \geq \lambda$, we obtain $l_n c_i \geq k_n c_{i+1}$ for all $n \geq 0$ and all $i \geq 0$, that is $c_{i+1} \leq \frac{l_n}{k_n} c_i$ and in the limit for $n \rightarrow \infty$, we get $c_{i+1} \leq \lambda c_i$ for all $i \geq 0$.

Similarly, consider a sequence of rational numbers $\frac{r_n}{s_n}$ converging to λ from the left. Taking $\alpha = r_n + s_n + 1$, $\beta = 1$, and $\gamma = s_n + 1$ in Lemma 3.3, and since the condition $(r_n + s_n + 1)S_0 + S_1 \leq (s_n + 1)(S_0 + S_1)$ is equivalent to $r_n S_0 \leq s_n S_1$ which holds

since $\frac{r_n}{s_n} \leq \lambda$, we obtain $r_n c_i \leq s_n c_{i+1}$ for all $n \geq 0$ and all $i \geq 0$, that is $c_{i+1} \geq \frac{r_n}{s_n} c_i$ and in the limit for $n \rightarrow \infty$, we get $c_{i+1} \geq \lambda c_i$ for all $i \geq 0$.

The two results imply that $c_{i+1} = \lambda c_i$ for all $i \geq 0$ where $0 \leq \lambda < 1$. Note that $\lambda \neq 1$ because $\sum_{i=0}^{\infty} c_i$ converges. \square

Since it is known that for $c_i = \lambda^i$, the weighted average payoff function induces memoryless optimal strategies in all two-player games, Theorem 3.1 shows that discounted sum is the only memoryless payoff function when the sum of weights $\sum_{i=0}^{\infty} c_i$ converges.

4 Weighted Average with Diverging Sum of Weights

In this section we consider weighted average objectives such that the sum of the weights $\sum_{i=0}^{\infty} c_i$ is divergent. We first consider the case when the sequence $(c_n)_{n \in \mathbb{N}}$ is bounded and show that the mean-payoff function is the only memoryless one.

4.1 Bounded sequence

We are interested in characterizing the class of weighted average objectives that are memoryless, under the assumption the sequence (c_n) is *bounded*, i.e., there exists a constant c such that $|c_n| \leq c$ for all n . The boundedness assumption is satisfied by the important special case of regular sequence of weights which can be produced by a deterministic finite automaton. We say that a sequence $\{c_n\}$ is *regular* if it is eventually periodic, i.e. there exist $n_0 \geq 0$ and $p > 0$ such that $c_{n+p} = c_n$ for all $n \geq n_0$. Recall that we assume the partial sum to be always non-zero, i.e., $d_n = \sum_{i=0}^{n-1} c_i \neq 0$ for all n . We show the following result.

Theorem 4.1. *Let $(c_n)_{n \in \mathbb{N}}$ be a sequence of real numbers with no zero partial sum such that $\sum_{i=0}^{\infty} |c_i| = \infty$ (the sum is divergent) and there exists a constant c such that $|c_i| \leq c$ for all $i \geq 0$ (the sequence is bounded). The weighted average payoff function ϕ defined by $(c_n)_{n \in \mathbb{N}}$ induces optimal memoryless strategies for all 2-player game graphs if and only if ϕ coincides with the mean-payoff function over regular words.*

Remark. From Theorem 4.1, it follows that all mean-payoff functions ϕ over bounded sequences that induce optimal memoryless strategies are equivalent to the mean-payoff function, in the sense that the optimal value and optimal strategies for ϕ are the same as for the mean-payoff function. This is because memoryless strategies induce a play that is a regular word. We also point out that it is not necessary that the sequence $(c_n)_{n \geq 0}$ consists of a constant value to define the mean-payoff function. For example, the payoff function defined by the sequence $c_n = 1 + 1/(n+1)^2$ also defines the mean-payoff function.

We prove Theorem 4.1 through a sequence of lemmas. In the following lemma we prove the existence of the limit of the sequence $\{\frac{1}{d_n}\}_{n \geq 0}$.

Lemma 4.1. *If $\liminf_{n \rightarrow \infty} \frac{1}{d_n} = 0$, then $\limsup_{n \rightarrow \infty} \frac{1}{d_n} = 0$.*

Proof. Since $l = \liminf_{n \rightarrow \infty} \frac{1}{d_n} = 0$, there is a subsequence $\{d_{n_k}\}$ which either diverges to $+\infty$ or $-\infty$.

1. If the subsequence $\{d_{n_k}\}$ diverges to $+\infty$, assume without loss of generality that each $d_{n_k} > 0$. Consider the one-player game graph G_3 shown in Figure 1. We consider the run corresponding to taking the edge with weight -1 for the first n_k steps followed by taking the 0 edge forever. The payoff for this run is given by

$$\liminf_{n \rightarrow \infty} \frac{-d_{n_k}}{d_n} = -d_{n_k} \cdot \limsup_{n \rightarrow \infty} \frac{1}{d_n} = -d_{n_k} \cdot L.$$

Since we assume existence of memoryless optimal strategies this payoff should lie between -1 and 0 . This implies that $d_{n_k} \cdot L \leq 1$ for all k . Since $L \geq l \geq 0$ and the sequence d_{n_k} is unbounded, we must have $L = 0$.

2. If the subsequence $\{d_{n_k}\}$ diverges to $-\infty$, assume that each $d_{n_k} < 0$. Consider the one-player game graph G_1 shown in Figure 1. We consider the run corresponding to taking the edge with weight 1 for the first n_k steps followed by taking the 0 edge forever. The payoff for this run is given by

$$\liminf_{n \rightarrow \infty} \frac{d_{n_k}}{d_n} = -|d_{n_k}| \cdot \limsup_{n \rightarrow \infty} \frac{1}{d_n} = -|d_{n_k}| \cdot L.$$

This payoff should lie between 0 and 1 (optimal strategies being memoryless), and this implies $L = 0$ as above. \square

Since $\limsup_{n \rightarrow \infty} d_n = \infty$, Lemma 4.1 concludes that the sequence $\{\frac{1}{d_n}\}$ converges to 0 i.e. $\lim_{n \rightarrow \infty} \frac{1}{d_n} = 0$. It also gives us the following corollaries which are a simple consequence of the fact that $\liminf_{n \rightarrow \infty} (a_n + b_n) = a + \liminf_{n \rightarrow \infty} b_n$ if a_n converges to a .

Corollary 4.1. *If $l = 0$, then the payoff function ϕ does not depend upon any finite prefix of the run, i.e., $\phi(a_1 a_2 \dots a_k u) = \phi(0^k u) = \phi(b_1 b_2 \dots b_k u)$ for all a_i 's and b_i 's.*

Corollary 4.2. *If $l = 0$, then the payoff function ϕ does not change by modifying finitely many values in the sequence $\{c_n\}_{n \geq 0}$.*

By Corollary 4.1, we have $\phi(xa^\omega) = a$ for all $a \in \mathbb{R}$. For $0 \leq i \leq k-1$, consider the payoff $S_{k,i} = \phi((0^i 1 0^{k-i-1})^\omega)$ for the infinite repetition of the finite sequence of k rewards in which all rewards are 0 except the $(i+1)$ th which is 1 . We show that $S_{k,i}$ is independent of i .

Lemma 4.2. *We have $S_{k,0} = S_{k,1} = \dots = S_{k,k-1} \leq \frac{1}{k}$.*

Proof. If $S_{k,0} \leq S_{k,1}$ then by prefixing by the single letter word 0 and using Lemma 2.1 we conclude that $S_{k,1} \leq S_{k,2}$. We continue this process until we get $S_{k,k-2} \leq S_{k,k-1}$. After applying this step again we get

$$S_{k,k-1} \leq \phi(0(0^{k-1} 1)^\omega) = \phi(1(0^{k-1} 1)^\omega) = \phi((10^{k-1})^\omega) = S_{k,0}.$$

Hence, we have $S_{k,0} \leq S_{k,1} \leq \dots \leq S_{k,k-1} \leq S_{k,0}$. Thus we have $S_{k,i}$ is a constant irrespective of the value of i . A similar argument works in the other case when $S_{k,0} \geq S_{k,1}$.




Fig. 2. The game $G(k, i)$.

Using the fact that $\liminf_{n \rightarrow \infty} (a_{1,n} + a_{2,n} + \dots + a_{k,n}) \geq \liminf_{n \rightarrow \infty} a_{1,n} + \dots + \liminf_{n \rightarrow \infty} a_{k,n}$, we get that $S_{k,i} \leq \frac{1}{k}$ for $0 \leq i \leq k-1$. \square

Let $T_{k,i} = -\phi((0^i(-1)0^{k-i-1})^\omega)$. By similar argument as in the proof of Lemma 4.2, we show that $T_{k,0} = T_{k,1} = \dots = T_{k,k-1} \geq \frac{1}{k}$.

We now show that (d_n) must eventually have always the same sign, i.e., there exists n_0 such that $\text{sign}(d_m) = \text{sign}(d_n)$ for all $m, n \geq n_0$. Note that by the assumption of non-zero partial sums, we have $d_n \neq 0$ for all n .

Lemma 4.3. *The d_n 's eventually have the same sign.*

Proof. Let $c > 0$ be such that $|c_n| < c$ for all n . Since (d_n) is unbounded, there exists n_0 such that $|d_n| > c$ for all $n > n_0$ and then if there exists $m > n_0$ such that $d_m > 0$ and $d_{m+1} < 0$, we must have $d_m > c$ and $d_{m+1} < -c$. Thus we have $c_{m+1} = d_{m+1} - d_m < -2c$, and hence $|c_{m+1}| > 2c$ which contradicts the boundedness assumption on (c_n) . \square

If the d_n 's are eventually negative then we use the sequence $\{c'_n = -c_n\}$ to obtain the same payoff and in this case $d_n = -\sum_{i=0}^{\infty} c'_i$ will be eventually positive. Therefore we assume that there is some n_0 such that $d_n > 0$ for all $n > n_0$. Let $\beta = \max\{|c_0|, |c_1|, \dots, |c_{n_0}|\}$. We replace c_0 by 1 and all c_i 's with β for $1 \leq i \leq n_0$. By corollary 4.2 we observe that the payoff function will still not change. Hence, we can also assume that $d_n > 0$ for all $n \geq 0$.

Lemma 4.4. *We have $S_{k,i} = \frac{1}{k} = T_{k,i}$ for all $0 \leq i \leq k-1$.*

Proof. Consider the game graph $G(k, i)$ which consists of state q_0 in which the player can choose among k cycles of length k where in the i th cycle, all rewards are 0 except on the $(i+1)$ th edge which has reward 1 (see Fig. 2).

Consider the strategy in state q_0 where the player after every $k \cdot r$ steps ($r \geq 0$) chooses the cycle which maximizes the contribution for the next k edges. Let i_r be the index such that $kr \leq i_r \leq kr + k - 1$ and $c_{i_r} = \max\{c_{kr}, \dots, c_{kr+k-1}\}$ for $r \geq 0$. The payoff for this strategy is $\liminf_{n \rightarrow \infty} t_n$ where $t_n = \frac{c_{i_0} + c_{i_1} + \dots + c_{i_{r-1}}}{d_n}$ for $i_{r-1} \leq n < i_r$.

Note that $c_{i_r} \geq \frac{\sum_{i=kr}^{kr+k-1} c_i}{k}$ (the maximum is greater than the average), and we get the following (where c is a bound on $(|c_n|)_{n \geq 0}$):

$$t_n \geq \frac{\sum_{i=0}^{n-1} c_i}{k \cdot d_n} - \frac{c}{d_n},$$

hence $\liminf_{n \rightarrow \infty} t_n \geq \frac{1}{k} - \liminf_{n \rightarrow \infty} \frac{c}{d_n} = \frac{1}{k}$.

By Lemma 4.2, the payoff of all memoryless strategies in $G(k, i)$ is $S_{k,0}$, and the fact that memoryless optimal strategies exist entails that $S_{k_0} = \liminf_{n \rightarrow \infty} t_n \geq \frac{1}{k}$, and thus $S_{k,0} = \frac{1}{k} = S_{k,i}$ for all $0 \leq i \leq k-1$.

Using a similar argument on the graph $G(k, i)$ with reward -1 instead of 1 , we obtain $T_{k,0} = \frac{1}{k} = T_{k,i}$ for all $0 \leq i \leq k-1$. \square

From Lemma 4.4, it follows that

$$S_{k,i} = \phi((0^i 1 0^{k-i-1})^\omega) = \lim_{n \rightarrow \infty} \frac{\sum_{r=0}^{\lfloor \frac{n}{k} \rfloor} c_{kr+i}}{d_n} = \frac{1}{k}$$

and hence,

$$\begin{aligned} \phi((a_0 a_1 \dots a_{k-1})^\omega) &= \liminf_{n \rightarrow \infty} \sum_{i=0}^{k-1} \left(a_i \cdot \frac{\sum_{r=0}^{\lfloor \frac{n}{k} \rfloor} c_{kr+i}}{d_n} \right) = \sum_{i=0}^{k-1} \left(a_i \cdot \lim_{n \rightarrow \infty} \frac{\sum_{r=0}^{\lfloor \frac{n}{k} \rfloor} c_{kr+i}}{d_n} \right) \\ &= \frac{\sum_{i=0}^{k-1} a_i}{k}. \end{aligned}$$

We show that the payoff of a regular word $u = b_1 b_2 \dots b_m (a_0 a_1 \dots a_{k-1})^\omega$ matches the mean-payoff value.

Lemma 4.5. *If $u := b_1 b_2 \dots b_m (a_0 a_1 \dots a_{k-1})^\omega$ and $v = (a_0 a_1 \dots a_{k-1})^\omega$ are two regular sequences of weights then $\phi(u) = \phi(v) = \frac{\sum_{i=0}^{k-1} a_i}{k}$.*

Proof. Let $r \in \mathbb{N}$ be such that $kr > m$. If $\phi(v) \leq \phi(0v)$ then using Lemma 2.1 we obtain $\phi(0v) \leq \phi(0^2v)$. Applying the lemma again and again, we get, $\phi(v) \leq \phi(0^m v) \leq \phi(0^{kr} v)$. From Corollary 4.1 we obtain $\phi(0^m v) = \phi(b_1 b_2 \dots b_m v) = \phi(u)$ and $\phi(0^{kr} v) = \phi((a_1 a_2 \dots a_k)^r v) = \phi(v)$. Therefore, $\phi(u) = \phi(v) = \frac{\sum_{i=0}^{k-1} a_i}{k}$. The same argument goes through for the case $\phi(v) \geq \phi(0v)$. \square

Proof (of Theorem 4.1). In Lemma 4.5 we have shown that the payoff function ϕ must match the mean-payoff function for regular words, if the sequence $\{c_n\}_{n \geq 0}$ is bounded. Since memoryless strategies in game graphs result in regular words over weights, it follows that the only payoff function that induces memoryless optimal strategies is the mean-payoff function which concludes the proof. \square

Observe that every regular sequence is bounded, and therefore the result of Theorem 4.1 holds for all weighted average objectives with divergent sum defined by regular sequence of weights.

Corollary 4.3. *Let $(c_n)_{n \in \mathbb{N}}$ be a regular sequence of real numbers with no zero partial sum such that $\sum_{i=0}^{\infty} |c_i| = \infty$ (the sum is divergent). The weighted average payoff function ϕ defined by $(c_n)_{n \in \mathbb{N}}$ induces optimal memoryless strategies for all two-player game graphs if and only if ϕ is the mean-payoff function.*

4.2 Unbounded sequence

The results of Section 3 and Section 4.1 can be summarized as follows: (1) if the sum of c_i 's is convergent, then the sequence $\{\lambda^i\}_{i \geq 0}$, with $\lambda < 1$ (discounted sum), is the only class of payoff functions that induce memoryless optimal strategies; and (2) if the sum is divergent but the sequence (c_n) is bounded, then the mean-payoff function is the only payoff function with memoryless optimal strategies (and the mean-payoff function is defined by the sequence $\{\lambda^i\}_{i \geq 0}$, with $\lambda = 1$). The remaining natural question is that if the sum is divergent and unbounded, then is the sequence $\{\lambda^i\}_{i \geq 0}$, with $\lambda > 1$, the only class that has memoryless optimal strategies. Below we show with an example that the class $\{\lambda^i\}$, with $\lambda > 1$, need not necessarily have memoryless optimal strategies.

We consider the payoff function given by the sequence $c_n = 2^n$. It is easy to verify that the sequence satisfies the partial non-zero assumption. We show that the payoff function does not result into memoryless optimal strategies. To see this, we observe that the payoff for a regular word $w = b_0 b_1 \dots b_t (a_0 a_1 \dots a_{k-1})^\omega$ is given by $\min_{0 \leq i \leq k-1} \left(\frac{a_i + 2a_{i+1} + \dots + 2^{k-1}a_{i+k-1}}{1+2+\dots+2^{k-1}} \right)$ i.e., the payoff for a regular word is the least possible weighted average payoff for its cycle considering all possible cyclic permutations of its indices (note that the addition in indices is performed modulo k).




Fig. 3. The game \mathcal{G}_{1024} .

Now, consider the game graph \mathcal{G}_{1024} shown in figure 3. The payoffs for both the memoryless strategies (choosing the left or the right edge in the start state) are $\min(\frac{5}{3}, \frac{4}{3})$ and $\min(\frac{4}{3}, \frac{8}{3})$ which are both equal to $\frac{4}{3}$. Although, if we consider the strategy which alternates between the two edges in the starting state then the payoff obtained is $\min(\frac{37}{15}, \frac{26}{15}, \frac{28}{15}, \frac{14}{15}) = \frac{14}{15}$ which is less than payoff for both the memoryless strategies. Hence, the player who minimizes the payoff does not have a memoryless optimal strategy in the game \mathcal{G}_{1024} . The example establishes that the sequence $\{2^n\}_{n \geq 0}$ does not induce optimal strategies.

Open question. Though weighted average objectives such that the sequence is divergent and unbounded may not be of the greatest practical relevance, it is an interesting theoretical question to characterize the subclass that induce memoryless strategies. Our counter-example shows that $\{\lambda^n\}_{n \geq 0}$ with $\lambda > 1$ is not in this subclass.

References

1. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.

2. M. Bojańczyk. Beyond omega-regular languages. In *Proc. of STACS*, LIPIcs 5, pages 11–16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2010.
3. A. Chakrabarti, L. de Alfaro, T. A. Henzinger, and M. Stoelinga. Resource interfaces. In *Proc. of EMSOFT*, LNCS 2855, pages 117–133. Springer, 2003.
4. K. Chatterjee and L. Doyen. Energy parity games. In *Proc. of ICALP: Automata, Languages and Programming (Part II)*, LNCS 6199, pages 599–610. Springer, 2010.
5. K. Chatterjee, L. Doyen, and T. A. Henzinger. Quantitative languages. *ACM Transactions on Computational Logic*, 11(4), 2010.
6. K. Chatterjee, T. A. Henzinger, B. Jobstmann, and R. Singh. Measuring and synthesizing systems in probabilistic environments. In *CAV 10*. Springer, 2010.
7. A. Church. Logic, arithmetic, and automata. In *Proceedings of the International Congress of Mathematicians*, pages 23–35. Institut Mittag-Leffler, 1962.
8. T. Colcombet and D. Niwiński. On the positional determinacy of edge-labeled games. *Theor. Comput. Sci.*, 352(1-3):190–196, 2006.
9. L. de Alfaro, T.A. Henzinger, and R. Majumdar. Discounting the future in systems theory. In *ICALP'03*, LNCS 2719, pages 1022–1037. Springer, 2003.
10. M. Droste and P. Gastin. Weighted automata and weighted logics. *Theor. Comput. Sci.*, 380(1-2), 2007.
11. A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *Int. Journal of Game Theory*, 8(2):109–113, 1979.
12. J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer-Verlag, 1997.
13. H. Gimbert and W. Zielonka. When can you play positionally? In *MFCS*, pages 686–697, 2004.
14. H. Gimbert and W. Zielonka. Games where you can play optimally without any memory. In *CONCUR'05*, pages 428–442. Springer, 2005.
15. Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games*, volume LNCS 2500. Springer, 2002.
16. V.A. Gurvich, A.V. Karzanov, and L.G. Khachiyan. Cyclic games and an algorithm to find minimax cycle means in directed graphs. *USSR Computational Mathematics and Mathematical Physics*, 28:85–91, 1988.
17. T. A. Henzinger. From boolean to quantitative notions of correctness. In *Proc. of POPL: Principles of Programming Languages*, pages 157–158. ACM, 2010.
18. A. Kechris. *Classical Descriptive Set Theory*. Springer, 1995.
19. E. Kopczyński. Half-positional determinacy of infinite games. In *Proc. of ICALP: Automata, Languages and Programming (2)*, LNCS 4052, pages 336–347. Springer, 2006.
20. T. A. Liggett and S. A. Lippman. Stochastic games with perfect information and time average payoff. *Siam Review*, 11:604–607, 1969.
21. D.A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.
22. J.F. Mertens and A. Neyman. Stochastic games. *International Journal of Game Theory*, 10:53–66, 1981.
23. A. Puri. *Theory of Hybrid Systems and Discrete Event Systems*. PhD thesis, University of California, Berkeley, 1995.
24. M.L. Puterman. *Markov Decision Processes*. John Wiley and Sons, 1994.
25. L.S. Shapley. Stochastic games. *Proc. Nat. Acad. Sci. USA*, 39:1095–1100, 1953.
26. W. Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, Beyond Words, chapter 7, pages 389–455. Springer, 1997.
27. Y. Velner and A. Rabinovich. Church synthesis problem for noisy input. In *Proc. of FOS-SACS*, LNCS 6604, pages 275–289. Springer, 2011.
28. U. Zwick and M. Paterson. The complexity of mean-payoff games on graphs. *Theoretical Computer Science*, 158:343–359, 1996.